

Prior Support Knowledge-Aided Sparse Bayesian Learning with Partly Erroneous Support Information

Jun Fang, Yanning Shen, Fuwei Li, and Hongbin Li, *Senior Member, IEEE*

Abstract—It has been shown both experimentally and theoretically that sparse signal recovery can be significantly improved given that part of the signal's support is known *a priori*. In practice, however, such prior knowledge is usually inaccurate and contains errors. Using such knowledge may result in severe performance degradation or even recovery failure. In this paper, we study the problem of sparse signal recovery when partial but partly erroneous prior knowledge of the signal's support is available. Based on the conventional sparse Bayesian learning framework, we propose a modified two-layer Gaussian-inverse Gamma hierarchical prior model and, moreover, an improved three-layer hierarchical prior model. The modified two-layer model employs an individual parameter b_i for each sparsity-controlling hyperparameter α_i , and has the ability to place non-sparsity-encouraging priors to those coefficients that are believed in the support set. The three-layer hierarchical model is built on the modified two-layer prior model, with a prior placed on the parameters $\{b_i\}$ in the third layer. Such a model enables to automatically learn the true support from partly erroneous information through learning the values of the parameters $\{b_i\}$. Variational Bayesian algorithms are developed based on the proposed hierarchical prior models. Numerical results are provided to illustrate the performance of the proposed algorithms.

Index Terms—Compressed sensing, sparse Bayesian learning, prior support knowledge.

I. INTRODUCTION

Compressed sensing is a recently emerged technique for signal sampling and data acquisition which enables to recover sparse signals from much fewer linear measurements [1]–[5]. In this paper, we study the problem of sparse signal recovery when prior information on the signal's partial support is available. In practice, prior information about the support region of the sparse signal may come from the support estimate obtained during a previous time instant. This is particularly the case for time-varying sparse signals whose support changes slowly over time. For example, in the real-time dynamic magnetic resonance imaging (MRI) reconstruction, it was shown that the support of a medical image sequence undergoes small variations with support changes (number of additions and removals) less than 2% of the support size [6]. Also, in the source localization problem, the locations of sources may

vary slowly over time. Thus the previous estimate of locations can be used as prior knowledge to enhance the accuracy of the current estimate.

The problem of sparse signal recovery with partial support information was studied in several independent and parallel works [6]–[9]. In [7], the partly known support is utilized to obtain a least-squares residual and compressed sensing is then performed on the least-squares residual instead of the original observation. When the partly known support is accurate, it is expected that the sparse signal associated with the residual has much fewer large nonzero components. Hence the least-squares residual-based compressed sensing can improve the recovery performance. Later in [6], a weighted ℓ_1 -minimization (modified basis pursuit) method was proposed, where the partially known support is excluded from the ℓ_1 -minimization (equivalent to setting the corresponding weights to zero). Sufficient conditions for exact reconstruction were derived, and it was shown that when a fairly accurate estimate of the true support is available, the exact reconstruction conditions are much weaker than the sufficient conditions for compressed sensing without utilizing the prior information [6]. This work was later extended to the noisy case [10], where a regularized modified basis pursuit denoising method was introduced and a computable bound on the reconstruction error was obtained. The weighted ℓ_1 -minimization approach was also studied in [8] by assuming a probabilistic support model which assigns a probability of being zero or nonzero to each entry of the sparse signal. The choice of the weights and the associated exact recovery conditions were investigated. In [9], a modified iterative reweighted method which incorporates the prior support information was proposed. Similar to [6], those weights corresponding to the *a priori* known support are assigned a very small value (close to zero), whereas other weights are recursively updated like conventional iterative reweighted methods. The problem of support knowledge-aided sparse signal recovery was also studied in a time-varying compressed sensing framework, e.g. [11], [12], where the temporal support correlation was modeled by a Markov chain [11] or a pattern-coupled structure [12]. Nevertheless, both works need to specify the inherent support correlation structure between two consecutive time steps, whereas our work here deals with a more general scenario involving no particular correlation structure between the prior support and the current support.

It has been observed [6]–[9] that the sparse recovery performance can be significantly improved through exploiting the prior support knowledge. Nevertheless, this improvement can only be achieved when the prior knowledge is fairly accu-

Jun Fang, Yanning Shen and Fuwei Li are with the National Key Laboratory of Science and Technology on Communications, University of Electronic Science and Technology of China, Chengdu 611731, China, Email: Jun-Fang@uestc.edu.cn

Hongbin Li is with the Department of Electrical and Computer Engineering, Stevens Institute of Technology, Hoboken, NJ 07030, USA, E-mail: Hongbin.Li@stevens.edu

This work was supported in part by the National Science Foundation of China under Grant 61172114, and the National Science Foundation under Grant ECCS-1408182.

rate. Existing methods, e.g. [6]–[9], may suffer from severe recovery performance degradation or even recovery failure in the presence of inaccurate prior knowledge. In practice, however, signal support estimation inevitably incurs errors, and in some cases, due to the support variation across time, the prior knowledge may contain a considerable amount of errors. In this paper, we first introduce a modified two-layer Gaussian-inverse Gamma prior model, where an individual parameter b_i , instead of a common parameter b used in the conventional framework, is employed to characterize each sparsity-controlling hyperparameter α_i . Through assigning different values to the parameters $\{b_i\}$, our new prior model has the flexibility to place non-sparsity-encouraging priors to those coefficients that are believed in the support set. Nevertheless, the above modified two-layer hierarchical model does not have a mechanism to learn the true support from the partly erroneous knowledge. To address this issue, we propose an improved hierarchical prior model which constitutes a three-layer hierarchical form, where a new layer is proposed in addition to the above modified two-layer hierarchical model. The new layer places a prior on the parameters $\{b_i\}$. Such an approach is capable of distinguishing the true support from erroneous support through learning the values of $\{b_i\}$. By resorting to the variational inference methodology, we develop a new sparse Bayesian learning method which has the ability to learn the true support from the erroneous information.

The rest of the paper is organized as follows. In Section II, we introduce a modified two-layer Gaussian-inverse Gamma hierarchical prior model and an improved three-layer hierarchical prior model which enables to learn the true support from partly erroneous knowledge. Variational Bayesian methods are developed in Section III. Simulation results are provided in Section IV, followed by concluding remarks in Section V.

II. HIERARCHICAL PRIOR MODEL

We consider the problem of recovering a sparse signal $\mathbf{x} \in \mathbb{R}^n$ from noise-corrupted measurements

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w} \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ ($m < n$) is the measurement matrix, and \mathbf{w} is the additive multivariate Gaussian noise with zero mean and covariance matrix $\sigma^2 \mathbf{I}$. Suppose we have partial but partly erroneous knowledge of the support of the sparse signal \mathbf{x} . The prior knowledge P can be divided into two parts: $P = S \cup E$, where S denotes the subset containing correct information and E denotes the error subset. If we let T denote the true support of \mathbf{x} and T^c denote the complement of the set T , i.e. $T \cup T^c = \{1, 2, \dots, n\}$, then we have $S \subset T$, and $E \subset T^c$. Note that the only prior information we have is P . The partition of S and E is unknown to us.

The prior support information can certainly be utilized to improve signal recovery. In the following, based on the conventional sparse Bayesian learning (SBL) framework, we first introduce a modified SBL hierarchical prior model which has the flexibility to place non-sparsity-encouraging priors to those coefficients that are believed in the support set. Furthermore, we propose an improved three-layer hierarchical

prior model which enables to learn the true support from the erroneous information and thus exploits the prior support information in a more constructive manner. To facilitate discussions and comparisons, we first provide a brief overview of the conventional SBL hierarchical model.

A. Overview of Conventional SBL

Sparse Bayesian learning was originally developed by Tipping in [13] to address regression and classification problems. Later on in [14]–[17], sparse Bayesian learning was adapted to solve the sparse recovery problem and obtained superior performance for sparse signal recovery in a series of experiments. In the conventional sparse Bayesian learning framework, a two-layer hierarchical prior model was proposed to promote the sparsity of the solution. In the first layer, \mathbf{x} is assigned a Gaussian prior distribution

$$p(\mathbf{x}|\boldsymbol{\alpha}) = \prod_{i=1}^n p(x_i|\alpha_i) \quad (2)$$

where $p(x_i|\alpha_i) = \mathcal{N}(x_i|0, \alpha_i^{-1})$, and $\boldsymbol{\alpha} \triangleq \{\alpha_i\}$, the inverse variance (precision) of the Gaussian distribution, are non-negative sparsity-controlling hyperparameters. The second layer specifies Gamma distributions as hyperpriors over the hyperparameters $\{\alpha_i\}$, i.e.

$$p(\boldsymbol{\alpha}) = \prod_{i=1}^n \text{Gamma}(\alpha_i|a, b) = \prod_{i=1}^n \Gamma(a)^{-1} b^a \alpha_i^{a-1} e^{-b\alpha_i} \quad (3)$$

where $\Gamma(a) = \int_0^\infty t^{a-1} e^{-t} dt$ is the Gamma function, the parameters a and b used to characterize the Gamma distribution are chosen to be very small values, e.g. 10^{-4} , in order to provide non-informative/uniform (over a logarithmic scale) hyperpriors over $\{\alpha_i\}$. As discussed in [13], a broad hyperprior allows the posterior mean of α_i to become arbitrarily large. As a consequence, the associated coefficient x_i will be driven to zero, thus yielding a sparse solution. This mechanism is also referred to as the “automatic relevance determination” mechanism which tends to switch off most of the coefficients that are deemed to be irrelevant, and only keep very few relevant coefficients to explain the data.

B. Modified Two-Layer Hierarchical Model

When the value of the parameter b is relatively large, e.g. $b = 0.5$, it can be readily observed from (3) that the hyperpriors are no longer uniform and now they encourage small values of $\{\alpha_i\}$. In this case, an arbitrarily large value of the posterior mean of α_i is prohibited. As a result, the two-layer hierarchical model no longer results in a sparsity-encouraging prior and therefore does not necessarily lead to a sparse solution. This fact inspires us to develop a new way to incorporate the prior support information into the sparse Bayesian learning framework. Specifically, instead of using a common parameter b for all sparsity-controlling hyperparameters $\{\alpha_i\}$, we hereby employ an individual parameter b_i for each hyperparameter α_i ,

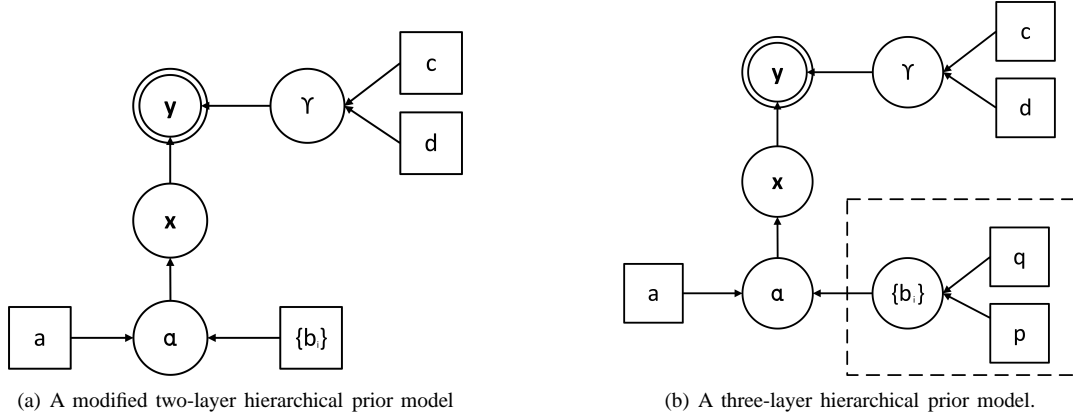


Fig. 1. Hierarchical models for support knowledge-aided sparse Bayesian learning.

i.e.

$$p(\alpha) = \prod_{i=1}^n \text{Gamma}(\alpha_i | a, b_i) = \prod_{i=1}^n \Gamma(a)^{-1} b_i^a \alpha_i^{a-1} e^{-b_i \alpha_i} \quad (4)$$

Such a formulation allows us to assign different priors to different coefficients. If the partial knowledge of the signal's support, P , is available, then the associated parameters of $\{b_i\}$ can be set to a relatively large value, say 0.5, in order to place a non-sparsity-encouraging prior on the corresponding coefficients, whereas the rest parameters of $\{b_i\}$ are still assigned a small value, say 10^{-4} , to encourage sparse coefficients, that is,

$$b_i = \begin{cases} 0.5 & i \in P \\ 10^{-4} & i \in P^c \end{cases} \quad (5)$$

where P^c denotes the complement of P , i.e. $P \cup P^c = \{1, 2, \dots, n\}$. The above modified hierarchical model seamlessly integrates the prior support information into the sparse Bayesian learning framework. Nevertheless, the modified two-layer hierarchical model which assigns fixed values to $\{b_i\}$ still lacks the flexibility to learn and adapt to the true situation. When the prior information contains a considerable portion of errors, this approach may suffer from significant performance loss and even recovery failure.

C. Proposed Three-Layer Hierarchical Model

To address the above issue, we partition the parameters $\{b_i\}$ into two subsets: $\{b_i, \forall i \in P\}$, and $\{b_i, \forall i \in P^c\}$. For $\{b_i, \forall i \in P^c\}$, the parameters are still considered to be deterministic and assigned a very small value, e.g. 10^{-4} ; while for $\{b_i, \forall i \in P\}$, instead of assigning a fixed large value, we model them as random parameters and place hyperpriors over these parameters. Since $\{b_i, \forall i \in P\}$ are positive values, suitable priors over $\{b_i, \forall i \in P\}$ are also Gamma distributions. In summary, we have

$$p(b_i) = \begin{cases} \text{Gamma}(b_i | p, q) = \Gamma(p)^{-1} q^p b_i^{p-1} e^{-q b_i} & i \in P \\ \delta(b_i - 10^{-4}) & i \in P^c \end{cases} \quad (6)$$

where $\delta(\cdot)$ denotes the Dirac delta function, p and q are parameters characterizing the Gamma distribution and their choice will be discussed in Section III. We see that the proposed model constitutes a three-layer hierarchical form which allows to learn the parameters $\{b_i, \forall i \in P\}$ in an automatic manner from the data. The proposed algorithm based on this prior model therefore has the ability to identify the correct support from erroneous information.

III. VARIATIONAL BAYESIAN INFERENCE

We now proceed to perform variational Bayesian inference for the proposed hierarchical models. Throughout this paper, the noise variance σ^2 is assumed unknown, and needs to be estimated along with other parameters. For notational convenience, define

$$\gamma \triangleq \sigma^{-2}$$

Following the conventional sparse Bayesian learning framework [13], we place a Gamma hyperprior over γ :

$$p(\gamma) = \text{Gamma}(\gamma | c, d) = \Gamma(c)^{-1} d^c \gamma^{c-1} e^{-d\gamma} \quad (7)$$

where the parameters c and d are set to small values, e.g. $c = d = 10^{-4}$. Before proceeding, we firstly provide a brief review of the variational Bayesian methodology.

A. Review of The Variational Bayesian Methodology

In a probabilistic model, let \mathbf{y} and $\boldsymbol{\theta}$ denote the observed data and the hidden variables, respectively. It is straightforward to show that the marginal probability of the observed data can be decomposed into two terms

$$\ln p(\mathbf{y}) = L(q) + \text{KL}(q || p) \quad (8)$$

where

$$L(q) = \int q(\boldsymbol{\theta}) \ln \frac{p(\mathbf{y}, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} d\boldsymbol{\theta} \quad (9)$$

and

$$\text{KL}(q || p) = - \int q(\boldsymbol{\theta}) \ln \frac{p(\boldsymbol{\theta} | \mathbf{y})}{q(\boldsymbol{\theta})} d\boldsymbol{\theta} \quad (10)$$

where $q(\boldsymbol{\theta})$ is any probability density function, $\text{KL}(q||p)$ is the Kullback-Leibler divergence between $p(\boldsymbol{\theta}|\mathbf{y})$ and $q(\boldsymbol{\theta})$. Since $\text{KL}(q||p) \geq 0$, it follows that $L(q)$ is a rigorous lower bound on $\ln p(\mathbf{y})$. Moreover, notice that the left hand side of (8) is independent of $q(\boldsymbol{\theta})$. Therefore maximizing $L(q)$ is equivalent to minimizing $\text{KL}(q||p)$, and thus the posterior distribution $p(\boldsymbol{\theta}|\mathbf{y})$ can be approximated by $q(\boldsymbol{\theta})$ through maximizing $L(q)$.

The significance of the above transformation is that it circumvents the difficulty of computing the posterior probability $p(\boldsymbol{\theta}|\mathbf{y})$ directly (which is usually computationally intractable). For a suitable choice for the distribution $q(\boldsymbol{\theta})$, the quantity $L(q)$ may be more amiable to compute. Specifically, we could assume some specific parameterized functional form for q and then maximize $L(q)$ with respect to the parameters of the distribution. A particular form of $q(\boldsymbol{\theta})$ that has been widely used with great success is the factorized form over the component variables $\{\theta_i\}$ in $\boldsymbol{\theta}$, i.e.

$$q(\boldsymbol{\theta}) = \prod_i q_i(\theta_i) \quad (11)$$

We therefore can compute the posterior distribution approximation by finding $q(\boldsymbol{\theta})$ of the form (11) that maximizes the lower bound $L(q)$. The maximization can be conducted in an alternating fashion for each latent variable, which leads to [18]

$$q_i(\theta_i) = \frac{\exp(\langle \ln p(\mathbf{y}, \boldsymbol{\theta}) \rangle_{k \neq i})}{\int \exp(\langle \ln p(\mathbf{y}, \boldsymbol{\theta}) \rangle_{k \neq i}) d\theta_i} \quad (12)$$

where $\langle \cdot \rangle_{k \neq i}$ denotes an expectation with respect to the distributions $q(\theta_i)$ for all $k \neq i$.

B. Bayesian Inference for Modified Two-Layer Model

Let $\boldsymbol{\theta} \triangleq \{\mathbf{x}, \boldsymbol{\alpha}, \gamma\}$ denote all hidden variables. We assume posterior independence among the variables \mathbf{x} , $\boldsymbol{\alpha}$, and γ , i.e.

$$\begin{aligned} p(\boldsymbol{\theta}|\mathbf{y}) &\approx q(\mathbf{x}, \boldsymbol{\alpha}, \gamma) \\ &= q_x(\mathbf{x}) q_\alpha(\boldsymbol{\alpha}) q_\gamma(\gamma) \end{aligned} \quad (13)$$

With this mean field approximation, the posterior distribution of each hidden variable can be computed by maximizing $L(q)$ while keeping other variables fixed using their most recent distributions, which gives

$$\begin{aligned} \ln q_x(\mathbf{x}) &= \langle \ln p(\mathbf{y}, \mathbf{x}, \boldsymbol{\alpha}, \gamma) \rangle_{q_\alpha(\boldsymbol{\alpha}) q_\gamma(\gamma)} + \text{constant} \\ \ln q_\alpha(\boldsymbol{\alpha}) &= \langle \ln p(\mathbf{y}, \mathbf{x}, \boldsymbol{\alpha}, \gamma) \rangle_{q_x(\mathbf{x}) q_\gamma(\gamma)} + \text{constant} \\ \ln q_\gamma(\gamma) &= \langle \ln p(\mathbf{y}, \mathbf{x}, \boldsymbol{\alpha}, \gamma) \rangle_{q_x(\mathbf{x}) q_\alpha(\boldsymbol{\alpha})} + \text{constant} \end{aligned}$$

where $\langle \cdot \rangle_{q(\cdot)}$ denotes the expectation with respect to (w.r.t.) the distribution $q(\cdot)$. In summary, the posterior distribution approximations are computed in an alternating fashion for each hidden variable, with other variables fixed. Details of this Bayesian inference scheme are provided below.

1). Update of $q_x(\mathbf{x})$: Keeping only the terms that depend on \mathbf{x} , the variational optimization of $q_x(\mathbf{x})$ yields

$$\begin{aligned} \ln q_x(\mathbf{x}) &\propto \langle \ln p(\mathbf{y}|\mathbf{x}, \gamma) + \ln p(\mathbf{x}|\boldsymbol{\alpha}) \rangle_{q_\alpha(\boldsymbol{\alpha}) q_\gamma(\gamma)} \\ &\propto -\frac{\langle \gamma \rangle}{2} (\mathbf{y} - \mathbf{A}\mathbf{x})^T (\mathbf{y} - \mathbf{A}\mathbf{x}) - \frac{1}{2} \mathbf{x}^T \langle \mathbf{D} \rangle \mathbf{x} \end{aligned} \quad (14)$$

where $\langle \gamma \rangle$ denotes the expectation w.r.t. $q_\gamma(\gamma)$, and $\langle \mathbf{D} \rangle \triangleq \text{diag}(\langle \alpha_1 \rangle, \dots, \langle \alpha_n \rangle)$, in which $\langle \alpha_i \rangle$ represents the expectation w.r.t. $q_\alpha(\boldsymbol{\alpha})$. It can be readily verified that $q_x(\mathbf{x})$ follows a Gaussian distribution with its mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Phi}$ given respectively as

$$\begin{aligned} \boldsymbol{\mu} &= \langle \gamma \rangle \boldsymbol{\Phi} \mathbf{A}^T \mathbf{y} \\ \boldsymbol{\Phi} &= \left(\langle \gamma \rangle \mathbf{A}^T \mathbf{A} + \langle \mathbf{D} \rangle \right)^{-1} \end{aligned} \quad (15)$$

2). Update of $q_\alpha(\boldsymbol{\alpha})$: Similarly, the approximate posterior $q_\alpha(\boldsymbol{\alpha})$ can be obtained by computing

$$\begin{aligned} \ln q_\alpha(\boldsymbol{\alpha}) &\propto \langle \ln p(\mathbf{x}|\boldsymbol{\alpha}) + \ln p(\boldsymbol{\alpha}|\mathbf{a}, \mathbf{b}) \rangle_{q_x(\mathbf{x})} \\ &\propto \sum_{i=1}^n \left\{ \left(a - \frac{1}{2} \right) \ln \alpha_i - \left(b_i + \frac{\langle x_i^2 \rangle}{2} \right) \alpha_i \right\} \end{aligned} \quad (16)$$

where $\langle x_i^2 \rangle$ denotes the expectation w.r.t. $q_x(\mathbf{x})$. Thus $\boldsymbol{\alpha}$ has a form of a product of Gamma distributions

$$q_\alpha(\boldsymbol{\alpha}) = \prod_{i=1}^n \text{Gamma}(\alpha_i; \tilde{a}, \tilde{b}_i) \quad (17)$$

in which the parameters \tilde{a} and \tilde{b}_i are respectively given as

$$\tilde{a} = a + \frac{1}{2} \quad \tilde{b}_i = b + \frac{1}{2} \langle x_i^2 \rangle \quad (18)$$

3). Update of $q_\gamma(\gamma)$: The approximate posterior distribution $q_\gamma(\gamma)$ can be computed as

$$\begin{aligned} \ln q_\gamma(\gamma) &\propto \langle \ln p(\mathbf{y}|\mathbf{x}, \gamma) + \ln p(\gamma|c, d) \rangle_{q_x(\mathbf{x})} \\ &\propto \left(\frac{m}{2} + c \right) \ln \gamma \\ &\quad - \left(\frac{1}{2} \langle (\mathbf{y} - \mathbf{A}\mathbf{x})^T (\mathbf{y} - \mathbf{A}\mathbf{x}) \rangle_{q_x(\mathbf{x})} + d \right) \gamma \end{aligned} \quad (19)$$

We can easily see that $q(\gamma)$ follows a Gamma distribution

$$q(\gamma) = \text{Gamma}(\gamma|\tilde{c}, \tilde{d}) \quad (20)$$

with the parameters \tilde{c} and \tilde{d} given respectively by

$$\begin{aligned} \tilde{c} &= \frac{m}{2} + c \\ \tilde{d} &= d + \frac{1}{2} \langle (\mathbf{y} - \mathbf{A}\mathbf{x})^T (\mathbf{y} - \mathbf{A}\mathbf{x}) \rangle_{q_x(\mathbf{x})} \end{aligned} \quad (21)$$

where

$$\langle (\mathbf{y} - \mathbf{A}\mathbf{x})^T (\mathbf{y} - \mathbf{A}\mathbf{x}) \rangle_{q_x(\mathbf{x})} = \|\mathbf{y} - \mathbf{A}\boldsymbol{\mu}\|_2^2 + \text{tr} \left\{ \mathbf{A}^T \mathbf{A} \boldsymbol{\Phi} \right\}$$

In summary, the variational Bayesian inference involves updates of the approximate posterior distributions for hidden variables \mathbf{x} , $\boldsymbol{\alpha}$, and γ . Some of the expectations and moments used during the update are summarized as

$$\begin{aligned} \langle \alpha_i \rangle &= \frac{\tilde{a}}{\tilde{b}_i} & \langle \gamma \rangle &= \frac{\tilde{c}}{\tilde{d}} \\ \langle x_i^2 \rangle &= \mu_i^2 + \phi_{i,i} \end{aligned}$$

where μ_i denotes the i th element of $\boldsymbol{\mu}$, and $\phi_{i,i}$ denotes the i th diagonal element of $\boldsymbol{\Phi}$. For clarity, we summarize our algorithm as follows.

Support Aided-SBL with No Support Learning

1. Given the current approximate posterior distributions $q_\alpha(\alpha)$ and $q_\gamma(\gamma)$, update the posterior distribution $q_x(\mathbf{x})$ according to (15).
2. Given $q_\gamma(\gamma)$ and $q_x(\mathbf{x})$, update $q_\alpha(\alpha)$ according to (18).
3. Given $q_\alpha(\alpha)$ and $q_x(\mathbf{x})$, update $q_\gamma(\gamma)$ according to (21).
4. Continue the above iterations until $\|\boldsymbol{\mu}^{(t+1)} - \boldsymbol{\mu}^{(t)}\|_2 \leq \epsilon$, where ϵ is a prescribed tolerance value.

The update for $\langle \alpha_i \rangle$ can be written as

$$\langle \alpha_i \rangle = \frac{1 + 2a}{\langle x_i^2 \rangle + 2b_i} \quad (22)$$

We observe that this update rule is similar to that of the conventional SBL, except that the common parameter b is now replaced by a set of individual parameters $\{b_i\}$. When b_i is very small, e.g. 10^{-4} , the update for $\langle \alpha_i \rangle$ is exactly the same as the conventional rule. Hence the Bayesian Occam's razor which contributes to the success of the conventional SBL also works here for those coefficients whose corresponding parameters $\{b_i\}$ are small. To see this, note that when computing the posterior mean and covariance matrix, a large $\langle \alpha_i \rangle$ tends to suppress the values of the corresponding components $\{\mu_i, \phi_{i,i}\}$ (c.f. (15)), which in turn leads to a larger $\langle \alpha_i \rangle$ (c.f. (22)). This negative feedback mechanism keeps decreasing most of the coefficients until they become negligible, while leaving only a few prominent nonzero entries survived to explain the data. On the other hand, when b_i is set to be relatively large, e.g. 0.5, from (22) we know that $\langle \alpha_i \rangle$ cannot be arbitrarily large and thus the automatic relevance determination mechanism is disabled. In other words, for those coefficients whose corresponding parameters $\{b_i\}$ are large, the priors assigned to these coefficients are no longer sparsity-encouraging. Hence the prior support information is seamlessly incorporated into the proposed algorithm by assigning different values to $\{b_i\}$.

C. Bayesian Inference for Proposed Three-Layer Model

When the prior support set P contains a considerable portion of errors, the above proposed algorithm may incur a significant performance loss because those parameters $\{b_i, \forall i \in E\}$ supposed to be very small are assigned large values. Instead, the proposed three-layer hierarchical model provides a flexible framework for adaptively learning values of $\{b_i, \forall i \in P\}$.

Let $\boldsymbol{\theta} \triangleq \{\mathbf{x}, \boldsymbol{\alpha}, \gamma, \bar{\mathbf{b}}\}$, where $\bar{\mathbf{b}} \triangleq \{b_i, \forall i \in P\}$ are hidden variables as well since they are assigned hyperpriors and need to be learned. We assume posterior independence among the hidden variables \mathbf{x} , $\boldsymbol{\alpha}$, γ , and $\bar{\mathbf{b}}$, i.e.

$$p(\mathbf{x}, \boldsymbol{\alpha}, \gamma, \bar{\mathbf{b}} | \mathbf{y}) \approx q(\mathbf{x}, \boldsymbol{\alpha}, \gamma, \bar{\mathbf{b}}) = q_x(\mathbf{x}) q_\alpha(\boldsymbol{\alpha}) q_\gamma(\gamma) q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}}) \quad (23)$$

With this mean field approximation, the posterior distribution of each hidden variable can be computed by minimizing the Kullback-Leibler (KL) divergence while keeping other

variables fixed using their most recent distributions, which gives

$$\begin{aligned} \ln q_x(\mathbf{x}) &= \langle \ln p(\mathbf{y}, \mathbf{x}, \boldsymbol{\alpha}, \gamma, \bar{\mathbf{b}}) \rangle_{q_\alpha(\boldsymbol{\alpha}) q_\gamma(\gamma) q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})} + \text{constant} \\ \ln q_\alpha(\boldsymbol{\alpha}) &= \langle \ln p(\mathbf{y}, \mathbf{x}, \boldsymbol{\alpha}, \gamma, \bar{\mathbf{b}}) \rangle_{q_x(\mathbf{x}) q_\gamma(\gamma) q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})} + \text{constant} \\ \ln q_\gamma(\gamma) &= \langle \ln p(\mathbf{y}, \mathbf{x}, \boldsymbol{\alpha}, \gamma, \bar{\mathbf{b}}) \rangle_{q_x(\mathbf{x}) q_\alpha(\boldsymbol{\alpha}) q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})} + \text{constant} \\ \ln q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}}) &= \langle \ln p(\mathbf{y}, \mathbf{x}, \boldsymbol{\alpha}, \gamma, \bar{\mathbf{b}}) \rangle_{q_x(\mathbf{x}) q_\alpha(\boldsymbol{\alpha}) q_\gamma(\gamma)} + \text{constant} \end{aligned}$$

Details of this Bayesian inference scheme are provided below.

1). Update of $q_x(\mathbf{x})$: The variational optimization of $q_x(\mathbf{x})$ can be calculated as follows by ignoring the terms that are independent of \mathbf{x} :

$$\begin{aligned} \ln q(\mathbf{x}) &\propto \langle \ln p(\mathbf{y} | \mathbf{x}, \gamma) + \ln p(\mathbf{x} | \boldsymbol{\alpha}) \rangle_{q_\alpha(\boldsymbol{\alpha}) q_\gamma(\gamma)} \\ &\propto -\frac{\langle \gamma \rangle}{2} (\mathbf{y} - \mathbf{A}\mathbf{x})^T (\mathbf{y} - \mathbf{A}\mathbf{x}) - \frac{1}{2} \mathbf{x}^T \langle \mathbf{D} \rangle \mathbf{x} \end{aligned} \quad (24)$$

We can easily verify that $q(\mathbf{x})$ follows a Gaussian distribution with its mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Phi}$ given respectively as

$$\begin{aligned} \boldsymbol{\mu} &= \langle \gamma \rangle \boldsymbol{\Phi} \mathbf{A}^T \mathbf{y} \\ \boldsymbol{\Phi} &= \left(\langle \gamma \rangle \mathbf{A}^T \mathbf{A} + \langle \mathbf{D} \rangle \right)^{-1} \end{aligned} \quad (25)$$

2). Update of $q_\alpha(\boldsymbol{\alpha})$: Similarly, the approximate posterior $q_\alpha(\boldsymbol{\alpha})$ can be computed as

$$\begin{aligned} \ln q_\alpha(\boldsymbol{\alpha}) &\propto \langle \ln p(\mathbf{x} | \boldsymbol{\alpha}) + \ln p(\boldsymbol{\alpha} | a, \mathbf{b}) \rangle_{q_x(\mathbf{x}) q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})} \\ &= \sum_i^n \langle (a - 0.5) \ln \alpha_i - (0.5x_i^2 + b_i) \alpha_i \rangle_{q_x(\mathbf{x}) q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})} \\ &\stackrel{(a)}{=} \sum_{i \in P} \{ (a + 0.5) \ln \alpha_i - (\langle b_i \rangle + 0.5 \langle x_i^2 \rangle) \alpha_i \} \\ &\quad + \sum_{i \in P^c} \{ (a + 0.5) \ln \alpha_i - (b_i + 0.5 \langle x_i^2 \rangle) \alpha_i \} \end{aligned} \quad (26)$$

where in (a), the terms inside the summation are partitioned into two subsets P and P^c because $\{b_i, i \in P^c\}$ are deterministic parameters, while $\{b_i, i \in P\}$ are latent variables and thus we need to perform the expectation over these hidden variables. The posterior $q(\boldsymbol{\alpha})$ has a form of a product of Gamma distributions

$$q(\boldsymbol{\alpha}) = \prod_{i=1}^n \text{Gamma}(\alpha_i | \tilde{a}, \tilde{b}_i) \quad (27)$$

with the parameters \tilde{a} and \tilde{b}_i given by

$$\tilde{a} = a + 0.5 \quad (28)$$

$$\tilde{b}_i = \begin{cases} \langle b_i \rangle + 0.5 \langle x_i^2 \rangle & i \in P \\ b_i + 0.5 \langle x_i^2 \rangle & i \in P^c \end{cases} \quad (29)$$

3). Update of $q_\gamma(\gamma)$: The approximate posterior $q_\gamma(\gamma)$ can be computed as

$$\begin{aligned} \ln q_\gamma(\gamma) &\propto \langle \ln p(\mathbf{y} | \mathbf{x}, \gamma) + \ln p(\gamma | c, d) \rangle_{q_x(\mathbf{x})} \\ &\propto \left(\frac{m}{2} + c - 1 \right) \ln \gamma \\ &\quad - \left(\frac{1}{2} \langle (\mathbf{y} - \mathbf{A}\mathbf{x})^T (\mathbf{y} - \mathbf{A}\mathbf{x}) \rangle_{q_x(\mathbf{x})} + d \right) \gamma \end{aligned} \quad (30)$$

It can be easily verified that $q(\gamma)$ follows a Gamma distribution

$$q(\gamma) = \text{Gamma}(\gamma|\tilde{c}, \tilde{d}) \quad (31)$$

where

$$\begin{aligned} \tilde{c} &= \frac{m}{2} + c \\ \tilde{d} &= d + \frac{1}{2} \langle (\mathbf{y} - \mathbf{A}\mathbf{x})^T (\mathbf{y} - \mathbf{A}\mathbf{x}) \rangle_{q_x(\mathbf{x})} \end{aligned} \quad (32)$$

in which

$$\langle (\mathbf{y} - \mathbf{A}\mathbf{x})^T (\mathbf{y} - \mathbf{A}\mathbf{x}) \rangle_{q_x(\mathbf{x})} = \|\mathbf{y} - \mathbf{A}\boldsymbol{\mu}\|_2^2 + \text{tr} \left\{ \mathbf{A}^T \mathbf{A} \boldsymbol{\Phi} \right\}$$

4). Update of $q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})$: The variational optimization of $q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})$ yields:

$$\begin{aligned} \ln q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}}) &\propto \langle \ln p(\boldsymbol{\alpha}|\mathbf{a}, \mathbf{b}) + \ln p(\bar{\mathbf{b}}|p, q) \rangle_{q_{\alpha}(\boldsymbol{\alpha})} \\ &\propto \sum_{i \in P} \{-b_i \langle \alpha_i \rangle + (p-1) \ln b_i - q b_i\} \end{aligned} \quad (33)$$

from which we can readily arrive at

$$q(\bar{\mathbf{b}}) = \prod_{i \in P} \text{Gamma}(b_i|p, \tilde{q}_i) \quad (34)$$

where

$$\tilde{q}_i = q + \langle \alpha_i \rangle$$

In summary, the variational Bayesian inference consists of successive updates of the approximate posterior distributions for hidden variables \mathbf{x} , $\boldsymbol{\alpha}$, γ , and $\bar{\mathbf{b}}$. Some of the expectations and moments used during the update are summarized as

$$\begin{aligned} \langle \alpha_i \rangle &= \frac{\tilde{a}}{\tilde{b}_i} & \langle \gamma \rangle &= \frac{\tilde{c}}{\tilde{d}} \\ \langle x_i^2 \rangle &= \mu_i^2 + \phi_{i,i} & \langle b_i \rangle &= \frac{p}{\tilde{q}_i} \end{aligned}$$

where μ_i denotes the i th element of $\boldsymbol{\mu}$, and $\phi_{i,i}$ denotes the i th diagonal element of $\boldsymbol{\Phi}$. We now summarize our algorithm as follows.

Partial Support Aided-SBL with Support Learning

1. Given the current approximate posterior distributions of $q_{\alpha}(\boldsymbol{\alpha})$, $q_{\gamma}(\gamma)$ and $q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})$, update $q_x(\mathbf{x})$ according to (25).
2. Given $q_x(\mathbf{x})$, $q_{\gamma}(\gamma)$ and $q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})$, update $q_{\alpha}(\boldsymbol{\alpha})$ according to (27)–(29).
3. Given $q_x(\mathbf{x})$, $q_{\alpha}(\boldsymbol{\alpha})$, and $q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})$, update $q_{\gamma}(\gamma)$ according to (31)–(32).
4. Given $q_x(\mathbf{x})$, $q_{\alpha}(\boldsymbol{\alpha})$ and $q_{\gamma}(\gamma)$, update $q_{\bar{\mathbf{b}}}(\bar{\mathbf{b}})$ according to (34).
5. Continue the above iterations until $\|\boldsymbol{\mu}^{(t)} - \boldsymbol{\mu}^{(t-1)}\|_2 \leq \epsilon$, where ϵ is a prescribed tolerance value. Choose $\hat{\boldsymbol{\mu}}^{(t)}$ as the estimate of the sparse signal.

The above proposed algorithm has the ability to adaptively learn the values of $\{b_i, \forall i \in P\}$, and thus has the potential to distinguish the correct support from erroneous information. To gain insight into the algorithm, we examine the update rules

for $\langle \alpha_i \rangle$ and $\langle b_i \rangle$, i.e. the posterior means of α_i and b_i . The update rules are given by

$$\langle \alpha_i \rangle = \begin{cases} \frac{1+2a}{\langle x_i^2 \rangle + 2\langle b_i \rangle} & i \in P \\ \frac{1+2a}{\langle x_i^2 \rangle + 2b_i} & i \in P^c \end{cases} \quad (35)$$

and

$$\langle b_i \rangle = \frac{p}{q + \langle \alpha_i \rangle} \quad (36)$$

respectively. We see that the two update rules are related to each other, with $\langle b_i \rangle$ used in updating $\langle \alpha_i \rangle$ and $\langle \alpha_i \rangle$ used in obtaining a new $\langle b_i \rangle$. A closer examination reveals that $\langle \alpha_i \rangle$ and $\langle b_i \rangle$ are inversely proportional to each other in their respective update rules. Specifically, a smaller $\langle b_i \rangle$ results in a larger $\langle \alpha_i \rangle$ (c.f. (35)), which in turn leads to a smaller $\langle b_i \rangle$ (c.f. (36)). This negative feedback mechanism could keep decreasing $\langle b_i \rangle$ until it becomes negligible, and eventually lead to an arbitrarily large $\langle \alpha_i \rangle$. Nevertheless, the interaction between $\langle b_i \rangle$ and $\langle \alpha_i \rangle$ could go the other way around, i.e. a larger $\langle b_i \rangle$ results in a smaller $\langle \alpha_i \rangle$, which in turn leads to a larger $\langle b_i \rangle$. In this case, $\langle \alpha_i \rangle$ will eventually converge to a finite value. The behavior of $\langle x_i^2 \rangle$ plays a key role in determining the course of the evolution, and in turn has an impact on the dynamic behavior of $\langle x_i^2 \rangle$ itself. With a proper choice of p and q , if $\langle x_i^2 \rangle$ becomes sufficiently small during the iterative process, then the process could evolve towards the $\langle \alpha_i \rangle \rightarrow \infty$ (that is, $\langle x_i^2 \rangle \rightarrow 0$) direction, otherwise the process will converge to a finite $\langle \alpha_i \rangle$, in which case a non-sparsity-encouraging prior is imposed on the coefficient x_i . This, clearly, is a sensible strategy to learn the parameters $\{b_i\}$. We now discuss the choice of the parameters p and q . To enable an efficient interaction between $\langle b_i \rangle$ and $\langle \alpha_i \rangle$, we hope $\langle \alpha_i \rangle$ plays a critical role in determining the value of $\langle b_i \rangle$. To this goal, the values of p and q should be set sufficiently small. Our experiments suggest that $p = q = 0.1$ is a suitable choice which enables effective support learning.

IV. SIMULATION RESULTS

We now carry out experiments to illustrate the performance of our proposed algorithms and compare with other existing methods. The proposed algorithm in Section III-B is referred to as the support knowledge-aided sparse Bayesian learning with no support learning (SA-SBL-NSL), and the one in Section III-C as the support knowledge-aided sparse Bayesian learning with support learning (SA-SBL-SL), respectively. The performance of the proposed algorithms¹ will be examined using both synthetic and real data. Throughout our experiments, the parameters p and q for our proposed algorithm are set equal to $p = q = 0.1$.

A. Synthetic Data

Suppose a K -sparse signal is randomly generated with the support set of the sparse signal randomly chosen according to a uniform distribution. The signals on the support set are independent and identically distributed (i.i.d.) Gaussian random

¹Matlab codes for our algorithm are available at <http://www.junfang-uestc.net/codes/SA-SBL.rar>

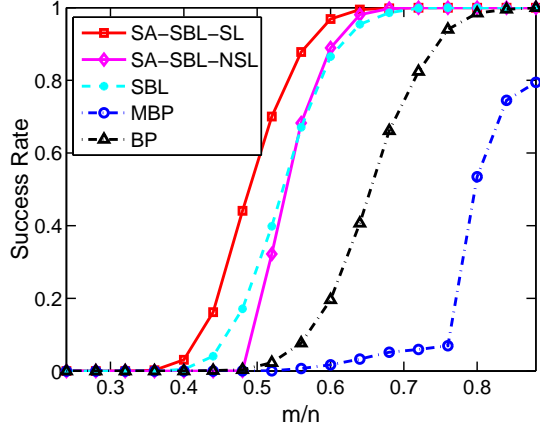


Fig. 2. Success rates of respective algorithms vs. the ratio m/n .

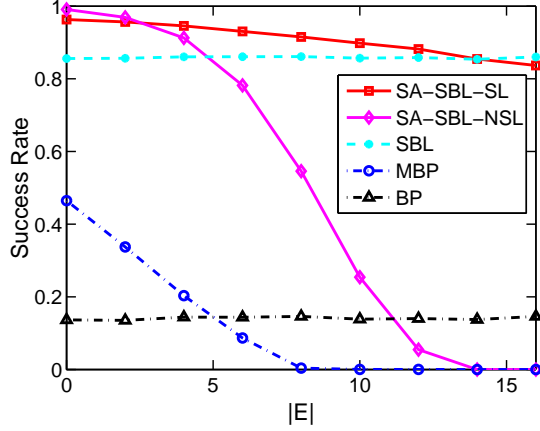


Fig. 3. Success rates of respective algorithms vs. the size of the error subset.

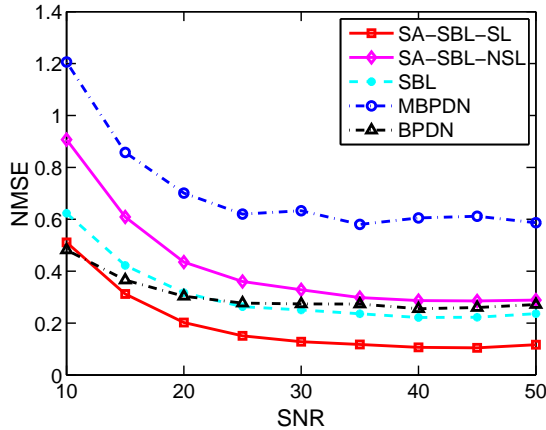


Fig. 4. Normalized mean-squared errors of respective algorithms vs. the signal-to-noise ratio.

variables with zero mean and unit variance. The measurement matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is randomly generated with each entry independently drawn from Gaussian distribution with zero mean and unit variance. The prior support information P consists of two subsets: $P = S \cup E$, where $S \subset T$ denotes

the subset containing the correct information, and $E \subset T^c$ is a subset comprised of false information. In our simulations, only the prior knowledge P is available, the exact partition of P into S and E is unknown. We compare our proposed algorithms with the conventional sparse Bayesian learning (SBL), the basis pursuit (BP) method, and the modified basis pursuit (MBP) method [6] which incorporates the partial support information by assigning different ℓ_1 -minimization weights to different coefficients.

We first consider the noiseless case. Fig. 2 plots the success rates of respective algorithms vs. the ratio m/n , where we set $K = 16$, $n = 50$, $|S| = 12$ and $|E| = 8$, $|S|$ and $|E|$ denote the cardinality (size) of the set S and E , respectively. The success rate is computed as the ratio of the number of successful trials to the total number of independent runs. A trial is considered successful if the normalized recovery error, i.e. $\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 / \|\mathbf{x}\|_2^2$, is no greater than 10^{-6} , where $\hat{\mathbf{x}}$ denotes the estimate of the true signal \mathbf{x} . Results are averaged over 1000 independent runs, with the measurement matrix and the sparse signal randomly generated for each run. It can be seen that our proposed SA-SBL-SL method presents a substantial performance advantage over the SA-SBL-NSL and the SBL methods. The performance gain is primarily due to the fact that the SA-SBL-SL method is able to learn the true support from the partly erroneous knowledge and thus make more effective use of the prior support information. We also observe that when a considerable number of errors are present in the prior knowledge, the methods SA-SBL-NSL and MBP present no advantage over their respective counterparts SBL and BP. To examine the behavior of the SA-SBL-SL method more thoroughly, we fix the number of elements in the set S and increase the number of elements in the error set E . Fig. 3 depicts the success rates vs. the number of elements in the error set E , where we set $m = 25$, $K = 16$, $|S| = 12$ and $|E|$ varies from 1 to 15. As can be seen from Fig. 3, when a fairly accurate knowledge is available, i.e. the number of errors is negligible or small, the SA-SBL-NSL achieves the best performance. This is an expected result since little learning is required at this point. Nevertheless, as the number of elements, $|E|$, increases, the SA-SBL-NSL suffers from substantial performance degradation. As compared with the SA-SBL-NSL, the SA-SBL-SL method provides stable recovery performance through learning the values of $\{b_i\}$, and outperforms all other algorithms when prior knowledge contains a considerable number of errors. We, however, notice that the proposed SA-SBL-SL method is surpassed by the conventional SBL method when inaccurate information becomes dominant (e.g. $|E| = 15$), in which case even learning brings limited benefits and simply ignoring the error-corrupted prior knowledge seems the best strategy.

We now consider the noisy case where the measurements are contaminated by additive noise. The observation noise is assumed multivariate Gaussian with zero mean and covariance matrix $\sigma^2 \mathbf{I}$. The normalized mean-squared errors (NMSEs) of respective algorithms as a function of signal-to-noise ratio (SNR) are plotted in Fig. 4, where we set $m = 25$, $n = 50$, $K = 16$, $|S| = 12$, and $|E| = 6$. The NMSE is calculated by averaging normalized squared errors over

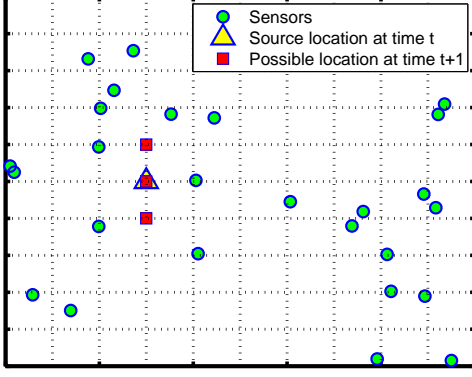


Fig. 5. A schematic diagram for source localization.

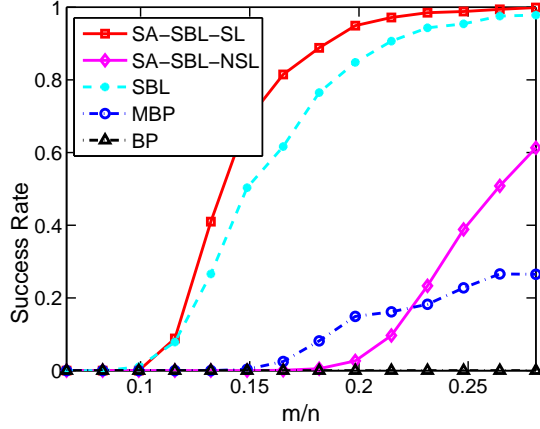


Fig. 6. Success rates of respective algorithms vs. the ratio m/n for the noiseless case.

10^3 independent runs. The SNR is defined as $\text{SNR}(\text{dB}) \triangleq 20 \log_{10}(\|\mathbf{Ax}\|_2/\|\mathbf{w}\|_2)$. The MBP-DN is a noisy version of the MBP method [10]. We observe that the conventional SBL and BP-DN methods outperform their respective counterparts: SA-SBL-NSL and MBP-DN. This, again, demonstrates that SA-SBL-NSL and MBP-DN methods are sensitive to prior knowledge inaccuracies. On the other hand, the proposed SA-SBL-SL method which takes advantage of the support learning in utilizing the prior knowledge, we assume the locations of these K_1 sources at time instant t are perfectly known to us and examine the recovery performance at time instant $t+1$. Fig. 6 depicts the success rates vs. the ratio m/n for the noiseless case. Results are averaged over 1000 independent runs, with sensors and locations of sources at time instant t randomly generated for each run. From Fig. 6, we see that the SA-SBL-NSL method yields performance much worse than the conventional SBL method. This is not surprising since the prior knowledge contains a substantial amount of erroneous information. It is also observed that the proposed SA-SBL-SL method which has the ability to learn the true support

B. Source Localization

We consider the problem of intensity-based source localization in sensor networks. The sensing field is partitioned into two-dimensional $n = 121$ -point virtual grid (Fig. 5) which is used to represent possible locations of the $K = 4$ sources. We have m randomly distributed sensors. The measurement collected at sensors can be written as

$$\mathbf{y}(t) = \mathbf{Ax}(t) + \mathbf{w}(t)$$

where $\mathbf{x}(t) \triangleq [x_1(t) \dots x_n(t)]^T$ is a sparse vector whose entry $x_i(t)$ denotes the intensity associated with the i th grid

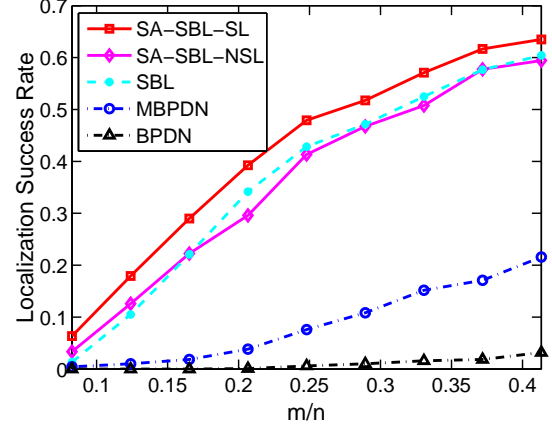


Fig. 7. Localization success rates of respective algorithms vs. the ratio m/n , SNR = 20dB.

point, $\mathbf{A} \triangleq [\mathbf{a}_1 \dots \mathbf{a}_m]^T$, $\mathbf{a}_i^T \triangleq [d_{1i}^{-\alpha} d_{2i}^{-\alpha} \dots d_{ni}^{-\alpha}]$, d_{ij} denotes the distance between the grid point i and the sensor j , and $\alpha = 2$ is the energy-decay factor.

In practice, sources may keep moving but the current locations' estimate could still be useful for the future localization. For example, we can expect that some sources may move to grid points close to their previous locations given that the interval between two time instants is sufficiently small. In our simulations, we assume that $K_1 = 3$ sources move slowly such that their next locations are partially predictable based on their current locations. Specifically, for these K_1 sources, each source either stays at its current position or moves to two of its immediate neighboring grid points² at the next time point (see Fig. 5). For simplicity, suppose $\hat{\mathcal{S}} = \{\hat{s}_1, \dots, \hat{s}_{K_1}\}$ is the set of locations associated with the K_1 sources at time instant t , then the set of possible locations of these K_1 sources at time instant $t+1$ is given by $P = \{\hat{s}_1 - 1, \hat{s}_1, \hat{s}_1 + 1, \dots, \hat{s}_{K_1} - 1, \hat{s}_{K_1}, \hat{s}_{K_1} + 1\}$. The set P can serve as prior support knowledge for source localization at time instant $t+1$. Nevertheless, the knowledge P is partly erroneous since all possible locations of these K_1 sources in the next move are included. For the rest $K - K_1$ sources, their locations are randomly chosen from the rest $n - |P|$ grid points. To test the effectiveness of respective algorithms in utilizing the prior knowledge, we assume the locations of these K_1 sources at time instant t are perfectly known to us and examine the recovery performance at time instant $t+1$. Fig. 6 depicts the success rates vs. the ratio m/n for the noiseless case. Results are averaged over 1000 independent runs, with sensors and locations of sources at time instant t randomly generated for each run. From Fig. 6, we see that the SA-SBL-NSL method yields performance much worse than the conventional SBL method. This is not surprising since the prior knowledge contains a substantial amount of erroneous information. It is also observed that the proposed SA-SBL-SL method which has the ability to learn the true support

²A grid point may have more than two neighboring points but we assume the source can only move to the specified two neighboring points, which is a reasonable assumption in practice since objects (e.g., vehicles) usually move along a pre-defined path such as a road or a highway.

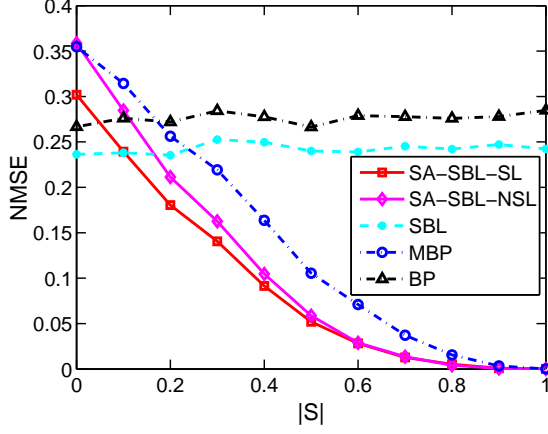


Fig. 8. Normalized mean squared errors of respective algorithms vs. the cardinality of the subset $|S|$.

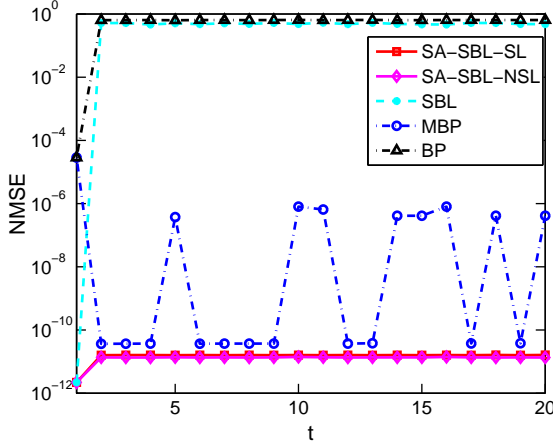


Fig. 9. Reconstruction of the sparsified 32×32 larynx sequence. Normalized mean squared errors vs. t .

from erroneous information achieves a significant performance improvement over the SA-SBL-NSL method. We now consider a noisy case where the measurements are corrupted by additive Gaussian noise. When noise is present, exact recovery of sparse signals is impossible. Nevertheless, accurate localization can still be achieved since reliable recovery of the support of sparse signals in the presence of noise is possible. In our simulations, locations of sources are estimated as the grid points associated with the largest K nonzero coefficients of the estimated signal. Fig. 7 plots the localization success rates as a function of the ratio m/n , where the signal-to-noise ratio (SNR) is set to 20dB. The localization success rate is calculated as the ratio of the number of successful trials to the total number of independent runs. A trial is considered successful if all K sources' locations are estimated correctly. Fig. 7, again, demonstrates the superiority of the proposed SA-SBL-SL method over the SA-SBL-NSL and the conventional SBL methods.

C. MRI Data

In this subsection, we carry out experiments on MRI images and sequences³. Images have sparse (or approximately sparse) structures in discrete wavelet transform (DWT) basis. By representing an image as a one-dimensional vector, the two-dimensional DWT (a two-level Daubechies-4 wavelet is used) of an image can be expressed as a product of an orthonormal matrix Ψ and the image vector x . The sensing matrix A is therefore equal to $A = Q\Psi$, where Q denotes the measurement acquisition matrix and its entries are i.i.d. normal random variables. We test all algorithms on sparsified images. Similar to [6], the image is sparsified by computing its 2D-DWT, retaining the coefficients from the 99%-energy support while setting others to zero and taking the inverse DWT.

We first evaluate the reconstruction performance of respective algorithms for a sparsified image. The image is a 32×32 larynx image obtained by resizing the 256×256 larynx image, i.e. $n = 1024$. For this image, its support size is $|T| = 70$. In our experiments, we fix the size of the set E to be $|E| = 0.4|T|$, and gradually increase the size of the set S . Fig. 8 depicts the NMSEs of respective algorithms vs. the size of S , where $m = 0.55n$ and $|S|$ varies from 0 to $0.9|T|$. Results are averaged over 500 independent runs, with the acquisition matrix Q , the sets E and S randomly generated for each run. From Fig. 8, we see that when the prior support knowledge is dominated by the error set E , the SBL and BP methods outperform their respective “support-aided” counterparts SA-SBL-NSL and MBP. Nevertheless, the SA-SBL-NSL and the MBP methods surpass and eventually achieve a significant performance improvement over the SBL and BP methods as the prior knowledge becomes more and more accurate. It can also be observed the SA-SBL-SL method presents uniform superiority over SA-SBL-NSL and MBP for different values of $|S|$, and the performance gap is wider in the small $|S|$'s region since learning in this region apparently brings more significant benefits as compared with learning in the large $|S|$'s region.

We also conduct a comparison for the sparsified 32×32 larynx sequence in Fig. 9. Since the support of the sequence undergoes a small variation, the prior support knowledge can be obtained as the estimate of the previous time instant. At the very beginning, i.e. $t = 0$, the prior knowledge set P is empty. Conventional SBL and BP methods are then used to obtain an initial support estimate for their respective support-aided methods. We set $m_0 = 0.6n$ in order to ensure a fairly accurate initial estimate is obtained. For $t > 0$, only $m_t = 0.3n$ measurements are collected for the signal reconstruction. From Fig. 9, we see that the all support-aided methods including SA-SBL-SL, SA-SBL-NSL and MBP are able to achieve exact reconstruction with only as few as $m = 0.3n$ measurements, whereas the conventional SBL and BP fail to recover the signal with these few measurements. Also, since the prior support knowledge is fairly accurate, support learning does not bring any additional benefits, and thus both the SA-SBL-SL and SA-SBL-NSL methods attain similar recovery performance.

³Available at http://www.phon.ox.ac.uk/jcoleman/Dynamic_MRI.html

V. CONCLUSIONS

We studied the problem of sparse signal recovery given that part of the signal's support is known *a priori*. The prior knowledge, however, may not be accurate and could contain erroneous information. This knowledge inaccuracy may result in considerable performance degradation or even recovery failure. To address this issue, we first introduced a modified two-layer Gaussian-inverse Gamma hierarchical prior model. The modified two-layer model employs an individual parameter b_i for each sparsity-controlling hyperparameter α_i , and therefore has the ability to place non-sparsity-encouraging priors to those coefficients that are believed in the support set. Based on this two-layer model, we then proposed an improved three-layer hierarchical prior model, with a prior placed on the parameters $\{b_i\}$ in the third layer. Such a model enables to automatically learn the true support from partly erroneous information through learning the values of $\{b_i\}$. Bayesian algorithms are developed by resorting to the mean field variational Bayes. Simulation results show that substantial performance improvement can be achieved through support learning since it allows us to make more effective use of the partly erroneous information.

REFERENCES

- [1] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.
- [2] E. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Information Theory*, no. 12, pp. 4203–4215, Dec. 2005.
- [3] D. L. Donoho, "Compressive sensing," *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, 2006.
- [4] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Information Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.
- [5] M. J. Wainwright, "Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting," *IEEE Trans. Information Theory*, vol. 55, no. 12, pp. 5728–5741, Dec. 2009.
- [6] N. Vaswani and W. Lu, "Modified-CS: modifying compressive sensing for problems with partially known support," *IEEE Trans. Signal Processing*, no. 9, pp. 4595–4607, September 2010.
- [7] N. Vaswani, "LS-CS-residual (LS-CS): compressive sensing on least squares residual," *IEEE Trans. Signal Processing*, no. 8, pp. 4108–4120, August 2010.
- [8] M. A. Khajehnejad, W. Xu, A. S. Avestimehr, and B. Hassibi, "Weighted ℓ_1 minimization for sparse recovery with prior information," in *Proc. IEEE Int. Symp. Information Theory*, Seoul, Korea, June 28–July 3 2009.
- [9] C. J. Miosso, R. von Borries, M. Argaez, L. Velazquez, C. Quintero, and C. M. Potes, "Compressive sensing reconstruction with prior information by iteratively reweighted least-squares," *IEEE Trans. Signal Processing*, no. 6, pp. 2424–2431, June 2009.
- [10] W. Lu and N. Vaswani, "Regularized modified BPDN for noisy sparse reconstruction with partial erroneous support and signal value knowledge," *IEEE Trans. Signal Processing*, no. 1, pp. 182–196, January 2012.
- [11] J. Ziniel and P. Schniter, "Dynamic compressive sensing of time-varying signals via approximate message passing," *IEEE Trans. Signal Processing*, no. 21, pp. 5270–5284, Nov. 2013.
- [12] J. Fang, Y. Shen, and H. Li, "Pattern coupled sparse bayesian learning for recovery of time-varying sparse signals," in *19th International Conference on Digital Signal Processing*, Hong Kong, August 20–23 2014.
- [13] M. Tipping, "Sparse Bayesian learning and the relevance vector machine," *Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.
- [14] D. P. Wipf, "Bayesian methods for finding sparse representations," Ph.D. dissertation, University of California, San Diego, 2006.
- [15] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Signal Processing*, vol. 56, no. 6, pp. 2346–2356, June 2008.
- [16] D. P. Wipf and B. D. Rao, "An empirical Bayesian strategy for solving the simultaneous sparse approximation problem," *IEEE Trans. Signal Processing*, vol. 55, no. 7, pp. 3704–3716, July 2007.
- [17] Z. Zhang and B. D. Rao, "Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 5, pp. 912–926, Sept. 2011.
- [18] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, "The variational approximation for Bayesian inference," *IEEE Signal Processing Magazine*, pp. 131–146, Nov. 2008.